

UNITED STATES PATENT APPLICATION

SERVICE SELECTION IN A SHARED ACCESS NETWORK

INVENTORS:

John W. Garrett

Charles Robert Kalmanek Jr.

Han Q. Nguyen

Kadangode K. Ramakrishnan

Cross Reference to Related Applications

This application claims priority to United States Provisional Application Serial No. 60/190,633, entitled "INTERNET SERVICE SELECTION OVER CABLE," filed on March 20, 2000, and to United States Provisional Application Serial No. 60/190,636, entitled "QUALITY OF SERVICE OVER THE HFC CABLE PLANT," filed on March 20, 2000, the contents of which are incorporated by reference herein.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25

SERVICE SELECTION IN A SHARED ACCESS NETWORK

Field of the Invention

5 The present invention relates generally to communication network services, and, more particularly, to providing multiple services in a communication network.

Background of the Invention

10 Customers of communication network services often desire access to a plurality of different services and different service providers. For example, when using a dial-up connection to a packet-switched data network such as the Internet, a customer can choose from multiple service providers by dialing different telephone numbers in the PSTN. The physical path from the customer to the customer's Internet Service Provider (ISP) is dedicated to the connection for
15 the duration of the telephone call. The ISP assigns an IP address to the customer and can link the authenticated customer and the assigned IP address to the physical address (e.g. dial-up modem) used by the customer. With this linkage, the ISP can ensure the customer only uses the address authorized by the ISP and can use the customer's IP address to manage access to the ISP's services. The
20 physical connection between a customer and the ISP, as well as the linkage to IP address assignment and customer authentication is terminated when the dial-up connection is terminated.

Constrained by the physical capacity of these temporary
connections across the PSTN, many service providers are moving to high-speed
25 access architectures (e.g., digital subscriber line (DSL), wireless, satellite, or cable) that provide dedicated physical connectivity directly to the subscriber and under the control of the ISP. These alternatives to shared access through the switched telephone network, however, do not lend themselves to shared access by multiple services and/or service providers.

Summary of the Invention

It is an object of the invention to enable multiple services or service providers to share the facilities of an access network infrastructure providing physical connectivity to subscribers. In accordance with an embodiment of the invention, bandwidth in the access network infrastructure is shared among the service networks. In accordance with another embodiment of the invention, the operator of the access network infrastructure uses hardware address bridging so that the service network providers can manage layer three operations, such as the assignment of network addresses.

These and other advantages of the invention will be apparent to those of ordinary skill in the art by reference to the following detailed description and the accompanying drawings.

Brief Description of the Drawings

FIG. 1 illustrates an interconnection of packet-switched service networks and an access network embodying principles of the invention.

FIG. 2 is a conceptual representation of an example embodiment of access network infrastructure sharing using layer one bandwidth allocation illustrating an aspect of the invention based on an HFC access architecture.

FIG. 3A and FIG. 3B is conceptual representation of an example embodiment using layer two hardware address bridging illustrating principles of the invention based on an HFC access architecture with corresponding end-to-end protocol layers.

FIG. 4A and FIG. 4B is conceptual representation of another example embodiment using layer two hardware address bridging illustrating principles of the invention based on an HFC access architecture with corresponding end-to-end protocol layers.

Detailed Description

In FIG. 1, a plurality of subscribers operating network access devices 101, 102, 103, ... 104 are provided access to communication network services, which are facilitated by a plurality of packet-switched data networks, shown in FIG. 1 as 151 and 152. Packet-switched data networks 151 and 152, referred to herein as "service networks," offer access to different services and/or are operated by different service providers. For example, service network 151 could provide packet-switched connectivity to public data networks while service network 152 could offer packet-switched telephony service (or the same public data network connectivity, but from a different service provider). The service networks, as is well known in the art, utilize a network addressing scheme to route datagrams to and from hosts: for example, where the service networks utilize the TCP/IP protocol suite, Internet Protocol (IP) addresses are assigned to each host and utilized in the process of routing packets from a source to a destination in the networks. See, e.g., "INTERNET PROTOCOL," IETF Network Working Group, RFC 791 (September 1981); S. Deering, R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification," IETF Network Working Group, RFC 1883 (December 1995), which are incorporated by reference herein. The invention shall be described herein with particular reference to the TCP/IP protocol suite and IP addresses, although those skilled in the art would readily be able to implement the invention using any of a number of different communication protocols.

The network access devices 101 ... 104 are typically customer premises equipment (CPE) such as a personal computer, information appliance, personal data assistant, data-enabled wireless handset, or any other type of device capable of accessing information through a packet-switched data network. Each network access device 101 ... 104 is either connected to or integrated with a network interface unit 111 ... 114, e.g. a modem, which enables communication through an access network infrastructure, shown as 120 in FIG. 1. The access network infrastructure 120 advantageously can be operated and maintained by an entity that is the same as or different from the entities operating and maintaining the service networks 151 and 152.

In a first aspect of the present invention, bandwidth in the access network infrastructure is shared among the various services/service providers. FIG. 2 sets forth an embodiment of this first aspect of the invention, described with particular reference to a hybrid fiber coaxial (HFC) access network. It should be recognized that those skilled in the art would readily be able to apply the principles of the present invention to other types of communication networks. As is known in the art, each network interface device 201 ... 202 is either connected to or integrated with a cable modem 211 which enables communication through the HFC network 221. In accordance with the Data Over Cable Service Interface Specification (DOCSIS), each service network 251 and 252 has a Cable Modem Termination System (CMTS), shown as 231 and 232 in FIG. 2A, which communicates with the cable modems 211 and manages access to both upstream and downstream cable capacity on the HFC networks 221. See, e.g., "Data-Over-Cable Service Interface Specifications: Cable Modem Termination System – Network Side Interface Specification," Cable Television Laboratories, Inc., SP-CMTS-NSI-I01-960702; "Data-Over-Cable Service Interface Specifications: Cable Modem to Customer Premise Equipment Interface Specification," Cable Television Laboratories, Inc., SP-CMCI-C02C-991015; "Data-Over-Cable Service Interface Specifications: Baseline Privacy Plus Interface Specifications," Cable Television Laboratories, Inc., SP-BPI+-I06-001215, which are incorporated by reference herein. The CMTS 231, 232 manages the scheduling of both upstream and downstream transmission and allocates cable capacity to individual customers identified by a Service ID (SID). The CMTS 231, 232 can have an integrated router or can be a separate device that bridges to a fast Ethernet switch which connects to a router. Connectivity is provided to the IP routers 241 and 242 in service networks 251 and 252, respectively. Accordingly, the HFC network 221 corresponds to the access network infrastructure 120 shown in FIG. 1.

Thus, the operator of the HFC network could allocate specific upstream and downstream channels to each IP service and/or service provider. Each service provider connects its own equipment (i.e. t, the CMTS) to the cable tuned to its allocated upstream and downstream channels. Customers of a service

would connect matching equipment (e.g., cable modems) to the cable tuned to the upstream and downstream channels allocated to the service. The service provider would allocate cable capacity within its allocated upstream and downstream
125 channels to customers using service provider specified procedures (e.g., DOCSIS). All protocols and procedures above RF transport are transparent to the operator of the HFC network.

The number of services/service providers that could provide service is constrained by the number of available upstream and downstream
130 channels. Channel allocation is static, so each service/service provider will require sufficient channel capacity to meet peak demand. Typically no more than eight to ten channels are available in the noisy 15 MHz to 50 MHz spectrum used for upstream channels, and a single service may require multiple upstream channels. For example, for telephony over cable (whether or not over IP) a single
135 1.6 MHz upstream channel supports no more than about 20 simultaneous phone calls (assuming non-compressed voice). Depending on cable deployment topology (e.g., number of homes passed), a successful telephony service may require several upstream cable channels to meet acceptable call blocking objectives, limiting the number of channels available for other services. The
140 quality of the upstream channels (i.e., noise) varies by frequency over time. That is, the quality of service would be different for different upstream channels (and therefore for different services/service providers), and would vary over time. To meet quality objectives, vendors are building CMTS and cable modems that can switch from one upstream channel to another based on real time channel quality
145 measurements. Allocation of cable channels to specific services/service providers limits this flexibility, reducing the quality of all two-way services (e.g., telephony, Internet Access) that share the same cable.

A cable modem is tuned to a single upstream channel and a single downstream channel. A customer desiring multiple services/service providers
150 would need a separate cable modem for each service/service provider. For example, a customer desiring simultaneous service from multiple ISPs (e.g., for parents and children) would need a separate cable modem for each ISP. Similarly,

a customer of both data and voice (over IP) service would need a separate cable modem for each service.

155 The service provider equipment (CMTS) connects directly to the HFC plant, and must be physically located at the head end of the cable. HUB/Head End floor space may limit the number of services/service providers that can be supported using this model. Deployment of service provider equipment in the cable operator HUB/Head End locations may complicate service
160 provider operations and maintenance, perhaps impacting customer service.

 The CMTS controls signal levels on the cable plant. A service provider with an improperly tuned CMTS could disrupt traffic to/from other service providers as well as other services on the shared cable.

 The service provider controls the physical connection to the
165 customer, including DOCSIS authentication, and can coordinate customer authentication and IP address assignment with DOCSIS authentication and policing. Services operate over different channels, ensuring each service/service provider only sees traffic of its own customers. All traffic is routed through the service provider giving the service provider control over quality of service.
170 Allocation of cable spectrum to services is unfortunately not usually an efficient use of scarce resources. The degenerate case of a single service provider and a single service uses scarce resources efficiently, but allocation of cable spectrum to multiple services and/or multiple service providers probably cannot scale beyond a few (3 to 5) services/service providers.

175 In a second aspect of the present invention, hardware address bridging can be used to and from a service provider's point-of-presence. The operator of the HFC network could bridge one or more DOCSIS LANs and service/service provider Points Of Presence (POPs). With this approach the cable operator maintains control of the cable, and the service provider manages IP
180 space. That is, the service provider assigns IP addresses and all packets are routed through the service provider. Depending on the scope of bridging done by the cable operator, service/service provider POPs are not necessarily in the same physical location as the CMTS. Bridging provides connectivity between

customers and service providers, but additional procedures are needed to isolate
 185 traffic to individual services/service providers. Without some form of traffic
 isolation, proprietary service provider information (e.g., customer MAC
 addresses) may be visible to all service providers. For example, a DHCP
 DISCOVER sent to the DOCSIS Broadcast address could be seen by every
 service provider on the bridged LAN. Moreover, any service provider's DHCP
 190 Server on the bridged LAN could respond to the DHCP DISCOVER, regardless
 of whose customer sent the packet.

Two exemplary embodiments of this approach to partitioning
 traffic are shown in FIG. 3 and 4. FIG. 3 shows the overlay of service/service
 provider specific Virtual LANs (VLANs) on the common bridged LAN, using
 195 procedures defined in the IEEE 802.1Q VLAN specification and potential
 enhancements. FIG. 4 shows use of point to point tunnels between customers and
 their selected service/service providers. The Point to Point Protocol is used to
 both encapsulate traffic and provide initialization and authentication procedures
 analogous to those used with dial-up access.

200 VLAN. The IEEE Virtual Bridged Local Area Networks
 specification provides protocols and procedures to bridge a collection of LAN
 segments, and partition the bridged LAN into multiple Virtual LANs (VLANs),
 each consisting of a subset of the bridged LAN segments. See IEEE 802.1Q/D11,
 205 IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged
 Local Area Networks, which is incorporated by reference herein. VLAN aware
 bridges add a four octet extension (including VLAN ID and priority) to the MAC
 header of frames that enter the VLAN aware bridged network, and delete the
 VLAN MAC header extension from frames that leave the VLAN aware bridged
 210 network. Unless other procedures are defined, the VLAN ID assigned to a frame
 is the VLAN ID assigned to the port through which the frame was received.
 However, the specification allows "proprietary extensions" of the specification to
 assign a VLAN ID to a frame based on information contained within the frame,
 such as MAC address or protocol ID. That is, a proprietary extension could

215 assign frames received from each customer on a DOCSIS LAN segment to a
VLAN associated with the customer's selected service/service provider, based on
a mapping of the customer's MAC address to the proper VLAN. Procedures
would be needed to establish the relationship between customer MAC address and
customer selected service/service provider. Procedures to dynamically change the
220 mapping of MAC address to service/service provider would be needed to enable
customers to access multiple services/service providers from the same device.

Consistent with normal bridging, a VLAN aware bridge transmits
frames through potential egress ports, potentially constrained by filters. Filters
may discard frames to prevent loops, or avoid unnecessary traffic (e.g., if the
225 destination MAC address is not reached through the port). A VLAN aware bridge
will also filter traffic based on a spanning tree associated with the VLAN
identified by the VLAN ID. That is, the VLAN aware bridge will not transmit the
frame through a port that is not part of the spanning tree associated with the
identified VLAN. In particular, a VLAN aware bridge will not transmit the frame
230 through a port to a (non-VLAN aware) LAN segment that is not included in the
identified VLAN.

Note that a single VLAN is identified in the VLAN MAC header
extension, perhaps based on the source MAC address of the frame (assuming a
"proprietary extension"). However, for egress, a port (and LAN segment) may be
235 included in multiple VLANs. A separate VLAN ID could be assigned to each
service/service provider, and the LAN Segment (e.g., LAN emulation over ATM,
Frame Relay, or Private Line) used to access the service/service provider. Using
administrative procedures, a customer's MAC address could be associated with
the VLAN ID of the customer selected service/service provider, and frames
240 originated by the customer could be assigned to the appropriate VLAN using a
"proprietary extension". The DOCSIS LAN segment could be assigned to each of
the VLANs of services/service providers selected by customers on the DOCSIS
LAN segment, allowing traffic from the customer selected service/service
provider to be transmitted over the customer's DOCSIS LAN.

245 The VLAN procedures described above would partition traffic so
that each service/service provider would receive frames only from its own
customers. A service/service provider could send traffic to customers of other
services/service providers (if it knew the customer's MAC addresses), since each
DOCSIS LAN would be a member of multiple VLANs. But VLAN procedures
250 would limit the scope of a customer's broadcast discovery procedures (e.g.,
DHCP DISCOVER) to the LAN segment of the customer selected service/service
provider (and perhaps other DOCSIS LAN segments).

If a LAN segment (e.g., a DOCSIS LAN segment) is a member of
multiple VLANs, procedures are needed to ensure multiple copies of an IP
255 multicast packet are not delivered to the LAN segment. If multiple customers on
a DOCSIS LAN segment are members of the same IP multicast group, and are
customers of multiple ISPs, each ISP may forward a copy of a the same multicast
packet to their own VLAN. Eventually, the CMTS that serves the DOCSIS LAN
would receive multiple copies of the same packet, all addressed to the multicast
260 address. At that point the CMTS cannot determine that the packets are copies of
the same packet without expensive processing of higher layer protocols, and
would likely forward each copy to the DOCSIS LAN. Multiple members of the
IP multicast group on the DOCSIS LAN would each receive multiple copies of
the same packet. Some higher layer protocols (e.g., TCP) should recognize
265 duplicate packets and recover, but other protocols (e.g., UDP) may be unaware
that the packets are duplicates. VLAN coupled with proprietary extensions can
provide a mechanism to partition services/service providers so that each
service/service provider would only receive traffic from its own customers.

270 PPPoE. Typically dial-up Internet access uses the point-to-point
protocol (PPP) to frame/transport IP packets. PPP over Ethernet (PPPoE)
provides a point to point connection over a (bridged) Ethernet LAN. IP packet
framing and associated PPP functionality is the same whether PPP is encapsulated
in Ethernet frames or encapsulated asynchronously over a dial-up connection.
275 From a service provider perspective, a PPPoE connection is functionally

equivalent to a PPP dial-up connection, enabling potential reuse of initialization and authentication procedures established for dial-up access.

280 A PPP connection over an Ethernet LAN is based on the MAC addresses of the end-points (i.e., PPP client and PPP server). This relationship is different than other PPP connections (e.g., dial-up connections) because the MAC address of a server is tied to a specific hardware component, leading to a potential single point of failure. (Conversely, a dial-up PPP connection is based on a telephone number that can be associated with a pool of modems and PPP servers.) PPPoE eliminates this potential single point of failure (server MAC address) by
285 leveraging LAN broadcast to discover the address of a PPP server. The PPPoE client sends a PPPoE Active Discovery Initiation (PADI) packet to the LAN broadcast address. Candidate PPP servers on the (bridged) LAN receive the PADI packet and may respond to the client.

290 PPP servers belonging to multiple service providers on the (bridged) LAN would all receive every PADI packet, and could respond, whether or not the service provider had a relationship with the customer. This potential exposure of service provider proprietary information and unauthorized access to customers could be problematic. This exposure can be eliminated if service provider PPP servers are not connected to the (bridged) LAN. Rather than deliver
295 PPP to service providers over Ethernet, PPP could be transported to each service provider encapsulated in UDP over IP (i.e., L2TP). Of course, without initialization procedures (such as provided by PPP), the customer does not have an IP address and can not communicate via IP. Instead, the cable operator could provide a type of relay service between customers and their selected service
300 providers. The cable operator could provide a Local Access Concentrator (LAC) to act as a "proxy" PPP server on the (bridged) LAN that would respond to customer PADI packets. Through the discovery procedure the cable operator LAC would learn the identity of the customer selected service/service provider. The cable operator LAC would relay PPP packets received from the customer (via
305 PPPoE) to the customer selected service provider via L2TP. In the opposite

direction the cable operator LAC would relay PPP packets received from the service provider (via L2TP) to the customer using PPPoE.

PPPoE could extend into a cable operator's network as far as the cable operator chooses to bridge the Ethernet (i.e., DOCSIS) LAN. The PPP connection could be extended to the service provider over an IP network using L2TP (perhaps over a point to point transport such as ATM, Frame Relay, or private lines). This hybrid preserves the PPP semantics and leverages the robustness of LAN discovery procedures without distributing proprietary customer information among multiple service providers. The service provider manages the customer IP space, including customer authentication (via PPP), assignment of IP addresses, and policing of IP address usage (linked to PPP tunnel). This solution may not use resources as efficiently as other solutions. L2TP adds considerable overhead to each packet, and does not leverage DOCSIS to deliver IP multicast efficiently. PPP framing provides the appearance of a point to point link to the IP layer. Current PPP specifications (e.g., PPPoE, RFC 2516) do not provide procedures to leverage underlying LAN multicast to transport IP multicast. Therefore, IP multicast requires PPP server replication for each customer subscribed to the multicast address, analogous to procedures used for dial-up PPP access.

The foregoing Detailed Description is to be understood as being in every respect illustrative and exemplary, but not restrictive, and the scope of the invention disclosed herein is not to be determined from the Detailed Description, but rather from the claims as interpreted according to the full breadth permitted by the patent laws. It is to be understood that the embodiments shown and described herein are only illustrative of the principles of the present invention and that various modifications may be implemented by those skilled in the art without departing from the scope and spirit of the invention. For example, the detailed description describes an embodiment of the invention with particular reference to an HFC access network architecture. However, the principles of the present invention could be readily extended to other access network architectures, such as

DSL, wireless, satellite, etc. Such an extension could be readily implemented by one of ordinary skill in the art given the above disclosure.